

Identification of stable models via nonparametric prediction error methods

Diego Romeres

Gianluigi Pillonetto

Alessandro Chiuso

Abstract—A new Bayesian approach to linear system identification has been proposed in a series of recent papers. The main idea is to frame linear system identification as predictor estimation in an infinite dimensional space, with the aid of regularization/Bayesian techniques. This approach guarantees the identification of stable predictors based on the prediction error minimization. Unluckily, the stability of the predictors does not guarantee the stability of the impulse response of the system. In this paper we propose and compare various techniques to address this issue. Simulations results comparing these techniques will be provided.

I. INTRODUCTION

Recent approaches for linear system identification describe the unknown system directly in terms of impulse response, thus describing an infinite dimensional model class. Needless to say, this is not entirely free of difficulties, since an alternative way to control the model complexity, i.e., to face the so called-bias variance tradeoff [1], [2], need to be found. It has been shown in the recent literature that the apparatus of Reproducing Kernel Hilbert Spaces (RHKS) or, equivalently, Bayesian Statistics provide powerful tools to face this tradeoff.

The paper [3] has shown how these infinite dimensional model classes can be used for identification of linear systems in the framework of prediction error methods, leading naturally to stable predictors. Yet stability of the predictor model does not necessarily guarantee stability of the so called “forward” (or simulation) model. As a matter of fact, we faced this stability issue when performing identification on a real data set from EEG recordings. Physical insight in this case suggests that the transfer function describing the link between potentials in different brain locations are expected to be stable, while the identified models where not.

Therefore, motivated by this real-world application, in this paper we shall tackle the problem of identifying stable (simulation) models when nonparametric prediction error methods [3] are used. We shall describe and compare, through an extensive simulation study, four possible solutions to this problem.

The paper is structured as follows: Section II formulates the problem. Sections III-V introduce four different approaches to guarantee stability of the identified models. Experimental results are described in Section VI and conclusions are drawn in Section VII.

Notation: Given a matrix M , M^\top shall denote its transpose, $\sigma(M)$ will be its eigenvalues. If $A(z)$ is a polynomial, $\sigma(A(z))$ will denote the set of roots of $A(z)$. Given two discrete time jointly stationary stochastic process $y(t)$ and $z(t)$, the symbol $\mathbb{E}[y(t)|z(s), s < t]$ shall denote the linear

minimum variance estimator (conditional expectation in the Gaussian case) of $y(t)$ given the past ($s < t$) history of $z(t)$.

II. STATEMENT OF THE MODEL STABILIZATION PROBLEM

We shall consider two jointly stationary discrete time zero mean stochastic processes $\{u(t)\}$, $\{y(t)\}$, $t \in \mathbb{Z}$, respectively the “input” and “output” processes.

As shown in [4], [5] under these assumptions there is an essentially unique representation of $y(t)$ in terms of $u(t)$ of the form

$$\begin{aligned} y(t) &= P(z)u(t) + H(z)e(t) \\ e(t) &:= y(t) - \mathbb{E}[y(t)|y(s), u(s), s < t] \end{aligned} \quad (1)$$

where

$$P(z) := \sum_{k=1}^{\infty} p_k z^{-k} \quad H(z) := \sum_{k=0}^{\infty} h_k z^{-k} \quad h_0 = 1 \quad (2)$$

and $H(z)$ is minimum-phase. This guarantees stability of the predictor $\hat{y}(t|t-1) := \mathbb{E}[y(t)|y(s), s < t]$:

$$\hat{y}(t|t-1) = H(z)^{-1} [(H(z) - 1)y(t) + P(z)u(t)] \quad (3)$$

In this paper we shall also assume that $P(z)$ (and thus $H(z)$) are stable¹ (i.e., analytic inside the open unit disc).

Prediction error approaches to system identification [1], [2] are based on estimating the predictor model

$$\begin{aligned} \hat{y}(t|t-1) &= F(z)y(t) + G(z)u(t) \\ F(z) &= \sum_{k=1}^{\infty} f_k z^{-k} \quad G(z) = \sum_{k=1}^{\infty} g_k z^{-k} \end{aligned} \quad (4)$$

Classic parametric methods [1], [2] start from a parametric description $P_\theta(z)$ and $H_\theta(z)$ of $P(z)$ and $H(z)$ in (1). This parametrization is usually constrained ($\theta \in \Theta$) so as to account for prior knowledge such as stability of $P_\theta(z)$, $H_\theta(z)$ and $H_\theta^{-1}(z)$. This induces a natural parametrization of the predictor $\hat{y}(t|t-1)$ which is thus denoted by $\hat{y}_\theta(t|t-1)$. Given a data set $y := \{y(t)\}_{t=1, \dots, T}$, $u := \{u(t)\}_{t=1, \dots, T}$, the parameters θ are then estimated minimizing the squared loss

$$\sum_{t=1}^T (y(t) - \hat{y}_\theta(t|t-1))^2. \quad (5)$$

More recently prediction error identification has been formulated in a nonparametric framework [3]. The main issue working in a nonparametric (possibly infinite dimensional) framework is that the problem of finding estimators \hat{f}, \hat{g} of $f := \{f_k\}_{k \in \mathbb{Z}^+}$ and $g := \{g_k\}_{k \in \mathbb{Z}^+}$ from measurements y, u is an ill-posed inverse problem [6]. The main idea, borrowed from [7] is to minimize the prediction error (5) searching for

¹Note that in a feedback configuration $P(z)$ is in principle allowed to be unstable provided there is a stabilizing feedback in action.

$\{f_k\}_{k \in \mathbb{Z}^+}$ and $\{g_k\}_{k \in \mathbb{Z}^+}$ in a suitable Reproducing Kernel Hilbert Space (RKHS) [8] which acts as a regularizer, also encoding notions of “stability” of the predictor (e.g. making sure that the estimated $\hat{F}(z)$ and $\hat{G}(z)$ are BIBO stable with probability one), see [3], [7] for details. Equivalently one can think that $\{f_k\}_{k \in \mathbb{Z}^+}$ and $\{g_k\}_{k \in \mathbb{Z}^+}$ are modeled as independent zero mean Gaussian Process [9] with a suitable covariance $K(t, s) = \text{cov}(f_t, f_s) = \text{cov}(g_t, g_s)$ (the same as the Reproducing Kernel above). This covariance is usually parametrized by some unknown hyperparameters η , which will be made explicit in the notation using a subscript, e.g. K_η and $p_\eta(f, g) = p_\eta(f)p_\eta(g)$. Under the assumption that the innovation process is Gaussian and independent of $f = \{f_k\}_{k \in \mathbb{Z}^+}$ and $g = \{g_k\}_{k \in \mathbb{Z}^+}$, also the marginal $p_\eta(y, u)$ and the posterior $p_\eta(f, g|y, u)$ are Gaussian, see [3] for details. The marginal density $p_\eta(y, u)$, also called marginal likelihood, can be used to estimate the unknown hyperparameter as:

$$\hat{\eta}_{ML} := \arg \max_{\eta} p_\eta(y, u). \quad (6)$$

Then, following the Empirical Bayes paradigm, estimators of $f = \{f_k\}_{k \in \mathbb{Z}^+}$ and $g = \{g_k\}_{k \in \mathbb{Z}^+}$ are then found from their posterior density $p_\eta(f, g|y, u)$ having fixed the hyperparameters to their estimated value $\hat{\eta}_{ML}$ [3]:

$$\hat{f} := \mathbb{E}_{\hat{\eta}_{ML}}[f|y, u], \quad \hat{g} := \mathbb{E}_{\hat{\eta}_{ML}}[g|y, u] \quad (7)$$

where $\mathbb{E}_{\hat{\eta}_{ML}}[\cdot]$ denotes conditional expectation having fixed $\eta = \hat{\eta}_{ML}$.

Unfortunately, BIBO stability of the impulse responses of $\{\hat{f}_k\}_{k \in \mathbb{Z}^+}$ and $\{\hat{g}_k\}_{k \in \mathbb{Z}^+}$ does not guarantee BIBO stability of the estimates

$$\hat{P}(z) := \frac{\hat{G}(z)}{1 - \hat{F}(z)}, \quad \hat{H}(z) := \frac{1}{1 - \hat{F}(z)} \quad (8)$$

of $P(z)$ and $H(z)$ in (1). In fact, BIBO stability of the sequences $\{\hat{f}_k\}_{k \in \mathbb{Z}^+}$ and $\{\hat{g}_k\}_{k \in \mathbb{Z}^+}$ have no relation with stability of $\hat{P}(z)$ and $\hat{H}(z)$ which, if no cancellations occur, depends on the zeros of $1 - \hat{F}(z) = 1 - \sum_{k=1}^{\infty} \hat{f}_k z^{-k}$.

For practical purposes when estimating the predictor model (4), the impulse responses $\{f_k\}_{k \in \mathbb{Z}^+}$ and $\{g_k\}_{k \in \mathbb{Z}^+}$ are truncated to a finite (yet arbitrarily large) p , so that we assume $F(z) = \sum_{k=1}^p f_k z^{-k}$, $G(z) = \sum_{k=1}^p g_k z^{-k}$.

Thus, the problem we consider in this paper, can be formulated as follows:

Problem 1: Given $y(t), u(t)$, $t \in [1, T]$, find $\{\hat{f}_k\}_{k \in [1, p]}$ and $\{\hat{g}_k\}_{k \in [1, p]}$ so that $\hat{P}(z)$ and $\hat{H}(z)$ in (8) are BIBO stable transfer functions. A sufficient generic² condition for this to happen is that

$$A(z) = z^p(1 - \sum_{k=1}^p \hat{f}_k z^{-k}) = z^p - [z^{p-1} \dots 1] \hat{f} \quad (9)$$

$$\hat{f} := [\hat{f}_1, \hat{f}_2, \dots, \hat{f}_p]^\top$$

is stable, i.e., has all roots inside $\mathcal{D} := \{z \in \mathbb{C} : |z| < 1\}$.

In the following we describe and compare three different techniques to achieve this aim. For each technique Problem 1

²i.e., if no cancellations occur, which is generic for estimated impulse responses.

is properly reformulated. In order to simplify the notation, in what follows, the input u will be dropped from the notation; therefore, for instance, we shall use $p_\eta(y)$ in lieu of $p_\eta(y, u)$.

III. STABILIZATION VIA LMI CONSTRAINT

The first stabilization technique is based on formulating stability of the model (8) as a constraint on the eigenvalues of the companion matrix of $A(z)$ in (9). This constraint can be characterized in terms of Linear Matrix Inequalities (LMI) as discussed in [10], and used later on in [11] to enforce stable models in subspace identification, thus leading to:

Problem 2 (Reformulation): Given a preliminary estimate $\tilde{f} := [\tilde{f}_1, \dots, \tilde{f}_p]^\top$, find a vector of coefficients \hat{f} so that

$$\hat{f} = \arg \min_{f \in \mathcal{F}_D} \|f - \tilde{f}\|^2 \quad (10)$$

where $\mathcal{F}_D := \{f \in \mathbb{R}^p : |\lambda| < 1 \forall \lambda \text{ s.t. } A(\lambda) = 0, A(z) = z^p - [z^{p-1} \dots 1] \hat{f}\}$, can be described by an LMI constraint as discussed below.

It should be observed that the use of the 2-norm in (10) is entirely arbitrary and, in fact, considering some form of model approximation error (e.g. difference of output predictors) would be preferable. In addition, when \tilde{f} is the outcome of a preliminary estimation step, a principled solution would require accounting for the distribution of \tilde{f} .

However, this brings in some technical difficulties related to the formulation of the quadratic problem, therefore, it is still subject of research.

Formulation of the LMI constraint

As shown in [10], a matrix F has all its eigenvalues in the LMI region $\mathcal{D} = \{z \in \mathbb{C} \text{ s.t. } f_{\mathcal{D}}(z) > 0\}$, where $f_{\mathcal{D}}(z)$ is an opportune polynomial matrix, if and only if there exists $P = P^\top \geq 0$ s.t.

$$M(F, P) = I_2 \otimes P + \left(\begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \otimes (FP) \right) + \left(\begin{matrix} * \\ * \end{matrix} \right)^\top \geq 0 \quad (11)$$

According to [11, Theorem 1], which presents small variations w.r.t the original central theorem in [10], we define the companion matrix of f as $\Psi(f) \in \mathbb{R}^{p \times p}$. Therefore, using (11), f is (Schur) stable if and only if $\exists P = P^\top \geq 0$ such that $M(\Psi(f), P) \geq 0$.

Unfortunately $M(\Psi(f), P)$ this is not linear in f and P since their product appears. Similarly to [11], this calls for a reparametrization of the constraint as follows: define the vector $\psi := Pf$ (so that $f = P^{-1}\psi$), $J := [0 \ I_{p-1}]$, and $M(\psi, P) := M(\Psi(f), P)$ i.e.,

$$M(\psi, P) = I_2 \otimes P + \left(\begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \otimes \begin{bmatrix} JP \\ \Psi^T \end{bmatrix} \right) + \left(\begin{matrix} * \\ * \end{matrix} \right)^\top \quad (12)$$

which is linear in ψ and P . Thus problem 2 can be reformulated as:

$$\hat{\psi}, \hat{P} = \arg \min_{\psi, P} \|\psi - P \hat{f}_B\|^2$$

$$\text{s. t. } M(\psi, P) \geq 0, \quad \text{Tr}(P) = p, \quad P = P^\top \geq 0 \quad (13)$$

where the constraint $Tr(P) = p$ is added to improve the numerical conditioning, see [11] for further details. The solution \hat{f} of Problem 2 is finally computed as:

$$\hat{f} = \hat{P}^{-1}\hat{\psi} \quad (14)$$

In the remaining of the paper the model $\hat{P}(z)$ obtained by plugging in (8) the estimators \hat{f} and \hat{g} obtained respectively from (14) and the Bayesian procedure in [3], will be called “LMI” model.

IV. STABILIZATION VIA PENALTY FUNCTION

The second stabilization technique is formulated to act directly inside the Bayesian procedure. As briefly discussed in section II, a crucial step of the Bayesian procedure is the estimation of the hyperparameter vector η through marginal likelihood optimization (6). It is in principle possible to restrict the set of admissible hyperparameters to a subset Ξ_S which lead to estimators (7) corresponding to stable models $\hat{P}(z)$ and $\hat{H}(z)$. This is not entirely trivial as the estimators (and thus the set Ξ_S) depend on the measured data y, u . This leads to the following:

Problem 3 (Reformulation): Estimate the hyperparameters η solving

$$\hat{\eta} = \arg \max_{\eta \in \Xi_S} p_{\eta}(y) = \arg \min_{\eta \in \Xi_S} -\ln p_{\eta}(y) \quad (15)$$

to the set $\Xi_S = \{\eta | A(z) \text{ Stable} \}$, i.e., the set of hyperparameters which lead to stable models $\hat{P}(z)$, $\hat{H}(z)$.

To force $\eta \in \Xi_S$, we can add a penalty function to the criterion in (15) which acts as a barrier to keep the estimate $\hat{\eta}$ away from the set of hyperparameters η leading to an unstable $A(z)$. In order to do so, we define $A_{\eta}(z)$ the polynomial $A(z)$ in (9) built with the estimator

$$\hat{f}_{\eta} := \mathbb{E}_{\eta}[f|y, u], \quad (16)$$

and $\bar{\rho}_{\eta} = \max |\sigma(A_{\eta}(z))|$. Next define the penalty function:

$$J(\bar{\rho}_{\eta}) = \frac{1}{(\alpha(\delta - \bar{\rho}_{\eta}))^{\alpha}} - \frac{1}{(\alpha\delta)^{\alpha}} \quad (17)$$

where $\delta \geq 1$ is a scalar which defines the barrier, α is a positive scalar which adjust how steep the barrier is.

As we can see in Figure 1 function (17) diverges ($J(\bar{\rho}_{\eta}) \rightarrow \infty$) when $\bar{\rho} \rightarrow \delta$ and $J(\bar{\rho}_{\eta}) \rightarrow 0$ when $\bar{\rho} \rightarrow 0$. Thus when (17) is added to the minimization problem (15), the solution $\bar{\rho}$ is pushed inside the stability region. The parameters α and δ are iteratively adjusted so as to guarantee that the final solution leads to a stable model, i.e. solves the constrained problem (15).

Notice that when $\alpha \rightarrow 0$, $J(\bar{\rho}_{\eta})$ gives no penalty for $\bar{\rho}_{\eta} < \delta$ and infinite penalty for $\bar{\rho}_{\eta} \geq \delta$. Elaborating upon the intuition above, it is easy to prove that the solution of Problem 3 can be found by the algorithm described below:

Algorithm 1:

1) Initialization:

- Compute η_0 using (15) and set $\alpha = 1$.
- Compute the predictor impulse response \hat{f}_{η_0} using (16), then determine the associated $A_{\eta_0}(z)$, $\bar{\rho}_{\eta_0}$.




Fig. 1. Representation of the penalty function $J(\bar{\rho}_{\eta})$. The red bullet represents the penalty function value associated to a specific $\bar{\rho}$ in an illustrative example. The blue (head filled) arrows show the effect of the penalty function on $\bar{\rho}$, the black (head no-filled) arrows the effects of changing the parameters α and δ . The blue dashed line represents the variation of $J(\bar{\rho}_{\eta})$ after a reduction of α .

2) While $\bar{\rho}_{\eta_k} \geq 1$

- Set $\delta = \bar{\rho}_{\eta_k}(1 + \epsilon)$
- Compute

$$\eta_k = \arg \min_{\eta} -\ln p_{\eta}(y) + J(\bar{\rho}_{\eta}) \quad (18)$$

and the associated $\bar{\rho}_{\eta_k}$

- If the value of $-\ln p_{\eta_k}(y) + J(\bar{\rho}_{\eta_k})$ is unchanged w.r.t. the $k - 1$ iteration, then perform the update: $\alpha = \alpha - \Delta\alpha, \delta = \delta - \Delta\delta$ where $\Delta\alpha$ and $\Delta\delta$ are chosen sufficiently small

3) Set $\alpha = \epsilon$ and $\delta = 1$.

Finally, the solution of Problem 3 is given by:

$$\hat{\eta} = \arg \min_{\eta} -\ln p_{\eta}(y) + J_p(\eta) \quad (19)$$

$$\hat{f} = \mathbb{E}_{\hat{\eta}}[f|y, u], \quad \hat{g} = \mathbb{E}_{\hat{\eta}}[g|y, u] \quad (20)$$

In the remaining of the paper the model obtained by (8) using (20) will be called “ML + PF” model.

Remark 1: Notice that the iterative procedure which updates δ and α is needed because, in general, it is not guaranteed that one can find an initial value of $\eta \in \Xi_S$. Note also that the set Ξ_S is always non-empty provided the hyperparameter vector η includes a scaling factor for the Kernel, i.e., a non negative scalar which multiplies the Kernel matrix. In fact, if this is the case, there exist values of η which lead to $\hat{f} = 0$ which, in turn leads to stable $\hat{P}(z)$ and $\hat{H}(z)$.

V. STABILIZATION VIA MARKOV CHAIN MONTE CARLO

In this Section we shall present a MCMC approach which yields the so called *full Bayes* estimator of f and g , introducing a (possibly non-informative) prior density ³ $p(\eta)$ on

³This may be a uniform distribution if the domain is compact.

the hyperparameter vector η . In order to enforce the stability constraint we consider the “stable” posterior distribution

$$p_S(f, g|y) = \frac{1}{p(y)} \int p(y|f, g) p_S(f, g|\eta) p(\eta) d\eta \quad (21)$$

where $p_S(f, g|\eta)$ is the “truncated” Gaussian prior

$$p_S(f, g|\eta) := \begin{cases} k_\eta p_\eta(f, g) & f : A(z) \text{ stable} \\ 0 & \text{otherwise} \end{cases} \quad (22)$$

which, a priori, excludes all impulse responses f which lead to unstable $A(z)$. Note that the constant k_η in (22) equals $k_\eta := \frac{1}{\int_{f \in \mathcal{F}} p(f, g|\eta) df dg}$, where $\mathcal{F} := \{f | A(z) \text{ stable}\}$. Unfortunately, the “stable” conditional

$$p_S(f, g|y, \eta) := \frac{p(y|f, g) p_S(f, g|\eta)}{p_S(y, \eta)}$$

is not Gaussian and, in addition, the integral in (21) cannot be computed in closed form. Therefore we tackle the problem using MCMC methods:

Problem 4 (Reformulation):

Obtain a sampling approximation of the “stable” posterior distribution (21). Compute from these samples the estimates \hat{f}, \hat{g} in (7) and \hat{P}, \hat{H} in (8) which satisfy the stability constraint. This will be done computing sample posterior means as well as sample MAP.

In order to sample from the stable posterior (21) one can use a Metropolis-Hasting type of algorithm as in [12].

We have now to address two fundamental issues for this algorithm to be implementable, namely:

- (i) Design the proposal density $Q_{f,g}(\cdot|\cdot)$
- (ii) Compute the posterior $p_S(f, g|y)$, up to a constant multiplicative factor⁴.

A preliminary step for both items (i) and (ii) is the computation of a set of samples $\eta_i \sim p(\eta|y)$ from the posterior of the hyperparameters, without accounting for the stability constraint.

In the next subsections we address these three issues.

Sampling from the posterior density $p(\eta|y)$

First, our aim is to draw points from the posterior density of η given y . Notice that:

$$p(\eta|y) = \frac{p_\eta(y) p(\eta)}{p(y)} \quad (23)$$

where, as mentioned earlier on, $p(\eta)$ is assumed to be a non informative prior distribution, and $p(y)$ is the normalization constant. The marginal density $p_\eta(y)$ of y given η can be computed in closed form, as discussed in [3] and is given by

$$p_\eta(y) = \exp\left(-\frac{1}{2} \ln(\det[2\pi\Sigma_\eta]) - \frac{1}{2} y^T \Sigma_\eta^{-1} y\right) \quad (24)$$

where

$$\Sigma_\eta = AK_\eta A^\top + BK_\eta B^\top + \sigma^2 I \quad (25)$$

where $\sigma^2 := \text{Var}\{e(t)\}$ is the variance of the innovation process (1) and A, B are matrices built with the past input-output data, see [3] for details.

⁴This is because only ratios of probabilities need to be computed.

In order to obtain samples from (23) we implemented a Metropolis-Hasting algorithm, see e.g. [12]. We are using a symmetric proposal distribution $q_\eta(\cdot|\cdot)$ which describes a random walk in the hyperparameter space, whose mean is centered in the present value and its variance contains information about the local curvature of the target. To do so, let us define:

$$\begin{aligned} \underline{\eta} &= \arg \min_{\eta} -\ln[p_\eta(y)p(\eta)] \\ H &= -\frac{d^2 \ln[p_\eta(y)p(\eta)]}{d\eta d\eta^\top} \end{aligned} \quad (26)$$

that is the Hessian matrix computed in $\underline{\eta}$. Thus we define $q_\eta(\cdot|\mu) = \mathcal{N}(\mu, \gamma H^{-1})$ where γ is a positive scalar chosen to obtain an acceptance probability in the MCMC algorithm around the 30% via a pilot analysis, see e.g. [13].

The acceptance rate of the MCMC results to be:

$$\alpha_{\eta_i} = \min\left(1, \frac{p_{\eta_i}(y)p(\eta_i)}{p_{\eta_{i-1}}(y)p(\eta_{i-1})}\right)$$

Proposal density

It is well known in the MCMC literature that an accurate choice of the proposal distribution may have a remarkable impact on the performance of the Markov Chain. In this paper we adopt a data-driven proposal computed from the posterior distribution disregarding the stability constraint. The algorithm we consider is based on the approximation

$$p(f, g|y) = \int_{\eta} p(f, g|y, \eta) p(\eta|y) d\eta \simeq \frac{1}{N} \sum_{i=1}^N p_{\eta_i}(f, g|y) \quad (27)$$

where $\eta_i, i = 1, \dots, N$ are the samples from $p(\eta|y)$ drawn by the MCMC algorithm above and

$$p_{\eta_i}(f, g|y) \sim \mathcal{N}(\mu_{\eta_i}^{MAP}, \Sigma_{\eta_i}^{MAP}) \quad (28)$$

is the (Gaussian) posterior density of f, g when the hyperparameters are fixed equal to η_i . The posterior means and variance are, respectively: $\mu_{\eta_i}^{MAP} := (\mathbb{E}_{\eta_i}[f|y], \mathbb{E}_{\eta_i}[g|y])$

$$\begin{aligned} \mathbb{E}_{\eta}[f|y] &= K_\eta A^\top \Sigma_\eta^{-1} y, \quad \mathbb{E}_{\eta}[g|y] = K_\eta B^\top \Sigma_\eta^{-1} y \\ \Sigma_\eta^{MAP} &= \underline{K}_\eta - \underline{K}_\eta \begin{bmatrix} A^\top \\ B^\top \end{bmatrix} \Sigma_\eta^{-1} \begin{bmatrix} A & B \end{bmatrix} \underline{K}_\eta \\ \underline{K}_\eta &= \begin{bmatrix} K_\eta & Q \\ Q & K_\eta \end{bmatrix} \end{aligned} \quad (29)$$

and Σ_η is defined in (25).

From (27) it follows that, in order to sample from the proposal density $p(f, g|y)$ one can

- 1) Sample $\eta_i \sim p(\eta|y)$
- 2) Sample $(f, g) \sim p_{\eta_i}(f, g|y)$ in (28)

Evaluation of the stable posterior $p_S(f, g|y)$

The stable posterior in equation (21) can be approximated as follows:

$$\begin{aligned} p_S(f, g|y) &= \int p_S(f, g, \eta|y) d\eta \\ &= \frac{1}{p(y)} \int p(y|f, g) p_S(f, g|\eta) p(\eta) \frac{q(\eta)}{q(\eta)} d\eta \\ &\simeq \frac{1}{N p(y)} \sum_{i=1}^N \frac{p(y|f, g) p_S(f, g|\eta_i) p(\eta_i)}{q(\eta_i)} \end{aligned} \quad (30)$$

with $\eta_i \sim q(\eta)$. Note that the quantities $p(y|f, g)$, $p_S(f, g|\eta)$ and $p(\eta)$ can be evaluated. Thus, setting $q(\eta) := p(\eta|y)$ and

using the MCMC algorithm described above to obtain samples from the posterior $p(\eta|y)$, the stable posterior $p_S(f, g|y)$ can then be approximated (up to the irrelevant normalization constant $p(y)$) from equation (30).

Algorithm

We are now ready to provide the MCMC algorithm to sample from the stable posterior $p_S(f, g|y)$ (21):

Algorithm 2 (MCMC):

Hyper-parameters MCMC:

- 1) Initialization: set $\eta_0 = \underline{\eta}$ using (26)
- 2) For $i > 0$ Iterate:
 - Sample η from $q_\eta(\cdot|\eta_{i-1}) \sim \mathcal{N}(\eta_{i-1}, \gamma H^{-1})$
 - Sample u from a uniform distribution on $[0, 1]$
 - Set $\eta_i = \begin{cases} \eta & \text{if } u \leq \frac{p_{\eta_i}(y)p(\eta_i)}{p_{\eta_{i-1}}(y)p(\eta_{i-1})} \\ \eta_{i-1} & \text{otherwise} \end{cases}$
- 3) After a burn-in period, keep the last N samples η_i which are (approximately) samples from $p(\eta|y)$.

Predictor Impulse Responses MCMC:

- 4) Initialization: compute $[f_0, g_0]$ from η_0 using (28)
 - 5) For $i = 1$ to N do
 - compute $\mu_{\eta_i}^{MAP}, \Sigma_{\eta_i}^{MAP}$ as in (29)
 - Sample (f', g') from $\mathcal{N}(\mu_i^{MAP}, \Sigma_i^{MAP})$
 - Compute α as
- $$\alpha := \min \left(1, \frac{p_S(f', g'|y)p(f^{(k)}, g^{(k)}|y)}{p_S(f^{(k)}, g^{(k)}|y)p(f', g'|y)} \right)$$
- with $p_S(f, g|y)$ and $p(f, g|y)$ approximated as in (30) and in (27).
- Sample u from a uniform distribution on $[0, 1]$
 - Set: $(f^{(i)}, g^{(i)}) = \begin{cases} (f', g') & \text{if } u \leq \alpha \\ (f^{(i-1)}, g^{(i-1)}) & \text{otherwise} \end{cases}$
- 6) The samples $(f^{(i)}, g^{(i)})$ obtained above are i.i.d. samples from $p_S(f, g|y)$ as requested by Problem 4. The estimates of $P(z)$ and $H(z)$ can be obtained as:
 - **Minimum Variance Estimate:** from each sample $(f^{(i)}, g^{(i)})$ compute the impulse responses $P_i(z)$ and $H_i(z)$ in (8) and compute the averages

$$\hat{P}(z) = \frac{1}{N} \sum_{i=1}^N P_i(z), \quad \hat{H}(z) = \frac{1}{N} \sum_{i=1}^N H_i(z) \quad (31)$$

We shall define $\hat{p} := \{\hat{p}_k\}_{k \in [1, p]}$, $\hat{h} := \{\hat{h}_k\}_{k \in [1, p]}$ the inverse \mathcal{Z} -transforms of \hat{P} and \hat{H} in (31).

- **Maximum a Posteriori Estimate**

$$\bar{f}, \bar{g} = \arg \max_{f, g} p_S(f, g|y) \quad (32)$$

In the remaining of the paper the model obtained by (8) using (31) and (32) will be called “MCMC posterior mean” model “MCMC MAP” model, respectively. Note that, from (31), an estimate of $P(z)$ is obtained directly. This is to guarantee that $P(z)$ is stable since the average $\sum_i \hat{P}_i(z)$ of BIBO stable function is BIBO stable. On the other hand, if one averaged⁵ the $f^{(i)}$ directly, there would be no guarantee

that the average f would lead to a stable $A(z)$ (and thus a stable model). Of course, if needed, an estimate of F can be obtained from (4) using \hat{P} and \hat{H} in (31) :

$$\hat{G}(z) := \hat{H}^{-1}(z)\hat{P}(z), \quad \hat{F}(z) := 1 - \hat{H}^{-1}(z)$$

VI. SIMULATIONS

The performance of the techniques described the paper are compared by means of a Monte Carlo experiment, considering identification or marginally stable models, i.e., with poles close to the complex unit circle. At each Monte Carlo run a 2^{nd} -order SISO ARMAX model, called M , is generated:

$$A(z)y(t) = kz^{-1}B(z)u(t) + C(z)e(t) \quad (33)$$

The two complex conjugate roots of the monic polynomial $A(z)$ are placed in $0.996 \cdot \exp(\pm j\frac{\pi}{3})$, $B(z)$ is a random polynomial whose roots are restricted to lie inside the circle of radius 0.9 and $C(z)$ has randomly roots chosen in the interval $[0.65, 0.73]$ so to ensure that the predictor impulse responses decay in no more than 30 steps.

The system input $u(t)$ and the disturbance noise $e(t)$ are independent white noise with unit variance (for both identification and test data sets). The constant k is designed so that the signal-to-noise ratio of (33) is one. More specifically, let $y_u(t) := B(z)/A(z)u(t)$ and $y_e(t) := C(z)/A(z)e(t)$, then k as been set to: $k = \sqrt{\text{var}(y_e)/\text{var}(y_u)}$. A Monte Carlo study of 5000 runs is implemented. At each run a model as (33) is used to generate an identification set of 400 samples and a test set of 1000 samples.

The predictor impulse responses f and g are estimated via the Bayesian System Identification described in [3] which is based on the Stable Spline Kernel as a priori covariance and the hyperparameters are determined as in (6). The predictor impulse responses are negligible for time lags larger than 30 and thus the truncation length is chosen as $p = 30$. The variance of the noise σ is computed via a low bias Least Square identification method. The estimators \hat{P} and \hat{H} in (8) obtained from the Stable Spline estimators \hat{f} , \hat{g} ended up being unstable about 150 times out of 5000 Monte Carlo runs. In these cases the stabilization procedures described in these paper have been applied. Thus our Monte Carlo analysis is limited to these 150 data sets which resulted in unstable systems.

The *CVX toolbox*, [15], which is based on *YALMIP*, was used in Matlab to solve the convex optimization problem (13), with solver *SeDuMi*, [16]. Instead, the Matlab function *fminsearch.m* has been used to solve problem (18).

Notice that all these unstable models have been stabilized by our techniques.

A. Performance results

In order to illustrate the identification performances, we first consider dominant poles of the estimated, which are shown in Figure 2; the horizontal line in 0.996 indicates the absolute value of the “true” dominant poles. All the estimation methods, and in particular “ML+PF” and “MCMC

⁵Recall that the average of stable polynomial is not necessarily a stable polynomial unless the degree is smaller than 3, see [14].




Fig. 2. Monte Carlo results. Boxplots of the absolute value of the dominant poles of the identified models. The horizontal line represents the absolute value of the dominant pole of the true model.

posterior mean”, tend to place the poles close to the unit circle.

In addition the estimated impulse responses are compared to the “true” ones in terms of relative errors on the estimated impulse responses:

$$\text{err}_i = \frac{1}{2} \frac{\|p_k - \hat{p}_k\|_2}{\|p_k\|_2} + \frac{1}{2} \frac{\|h_k - \hat{h}_k\|_2}{\|h_k\|_2} \quad (34)$$

where $\{\hat{p}\}$ and $\{\hat{h}\}$ are the estimators of the true impulsed responses $\{p\}$ and $\{h\}$.




Fig. 3. Monte Carlo results. Boxplots of the $\{\text{err}_i\}$.

Figure 3 reports the Boxplots of $\{\text{err}_i\}$ for the estimated models. The “MCMC posterior mean” estimator outperforms all the others significantly. The remaining three techniques yield rather poor quality properties in the identification of the system which is due to poor estimation of “dominant” modes. Indeed, a higher absolute value of the dominant pole corresponds to a slower decay rate of the impulse responses. When the estimators place a dominant pole very close to the unit circle, this results in a significant tail in the impulse response which, in turn yield a very high relative error.

The algorithm “MCMC posterior mean” deserves a separate discussion. In this case, since the estimated $P(z)$ is the average of all $\hat{P}_i(z)$ s, the dominant pole of \hat{P} is the slowest among the dominant poles of \hat{P}_i s. Yet, the effect of these dominant modes on the relative error is mitigated by the factor $\frac{1}{N}$ in the average (31).

VII. CONCLUSION

We have presented four different techniques to face the problem of identifying a stable system using a Bayesian framework based on the minimization of the predictor error. The experiment shows all methods ultimately produce stable models which perform comparably in terms of prediction error (not reported for reasons of space); however, only the model estimated with the so called “MCMC posterior mean” technique perform satisfactorily in terms of impulse response fit. In future work, we will discuss new techniques to overcome the problem in the identification performance without the usage of a MCMC. In particular, we are looking for new regularization which take in account penalty term both in the predictor and in the system impulse responses.

ACKNOWLEDGEMENT

This work has been supported by MIUR through the FIRB project “Learning meets time” (RBFR12M3AC)

REFERENCES

- [1] L. Ljung, *System Identification, Theory for the User*. Prentice Hall, 1997.
- [2] T. Söderström and P. Stoica, *System Identification*. Prentice-Hall, 1989.
- [3] G. Pillonetto, A. Chiuso, and G. D. Nicolao, “Prediction error identification of linear systems: a nonparametric gaussian regression approach,” *Automatica*, no. 47, pp. 291–305, 2011.
- [4] T. Söderström, L. Ljung, and I. Gustafsson, “Identifiability conditions for linear multivariable systems operating under feedback,” *IEEE Trans. on Aut. Contr.*, vol. 21, pp. 837–840, 1976.
- [5] T. Ng, G. Goodwin, and B. Andersson, “Identifiability of mimo linear dynamic systems operating in closed loop,” *Automatica*, vol. 13, pp. 477–485, 1977.
- [6] A. Tikhonov and V. Arsenin, *Solutions of Ill-Posed Problems*. Washington, D.C.: Winston/Wiley, 1977.
- [7] G. Pillonetto and G. De Nicolao, “A new kernel-based approach for linear system identification,” *Automatica*, vol. 46, no. 1, pp. 81–93, 2010.
- [8] N. Aronszajn, “Theory of reproducing kernels,” *Trans. of the American Mathematical Society*, vol. 68, pp. 337–404, 1950.
- [9] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- [10] M. Chilali and P. Gahinet, “H-infinity design with pole placement constraints: an lmi approach,” *IEEE Transactions on Automatic Control*, vol. 41, pp. 358–367, 1996.
- [11] D. N. Miller and R. A. de Callafon, “Subspace identification with eigenvalue constraints,” *Automatica*, vol. 49, pp. 2468–2473, 2013.
- [12] S. R. W. Gilks and D. Spiegelhalter, *Markov Chain Monte Carlo in Practice*. London: Chapman and Hall, 1996.
- [13] G. Roberts, A. Gelman, and W. Gilks, *Weak convergence and optimal scaling of random walk metropolis algorithms*. Ann. Appl. Prob., 1997, vol. 7.
- [14] M. Gora, “Stability of the convex combination of polynomials,” *Control and Cybernetics*, 2007.
- [15] M. Grant, S. Boyd, and Y. Ye, “Disciplined convex programming,” in *Global Optimization: from Theory to Implementation, Nonconvex Optimization and Its Applications*, L. Liberti and N. Maculan, Eds. New York: Springer, 2006, pp. 155–210.
- [16] J. Sturm, “Using sedumi, a matlab toolbox for optimization over symmetric cones,” 2001.